170-WP-010-001

ECS Capabilities Supporting a Federated System

White Paper--Not intended for formal review or government approval.

August 1996

Prepared Under Contract NAS5-60000

RESPONSIBLE ENGINEERS

Cheryl Croft /s/ for — 9/3/96
Jerry Pisarcik, Senior Scientist/Engineer — Date
EOSDIS Core System Project

Cheryl Croft /s/ — 9/3/96
Cheryl Croft, Principal Scientist/Engineer — Date
EOSDIS Core System Project

Paul Krauss /s/ — 9/3/96
Paul Krauss, Systems Engineer — Date
EOSDIS Core System Project

SUBMITTED BY

Paul Fingerman /s/ — 9/3/96
Paul Fingerman, Office Manager — Date
EOSDIS Core System Project

Hughes Information Technology Systems
Upper Marlboro, Maryland

This page intentionally left blank.

Abstract

The Earth Observing System Data and Information System (EOSDIS) is evolving toward an increasingly distributed federation of information providers and users. Through this federated approach, data production algorithms, analysis tools, and computing resources should become more accessible to the user community. And as the number of global change information providers grows, the diversity of information and services available to the earth science research community should increase. At the center of the federated approach is the Earth Science Information Partner, or ESIP. Within the federation, ESIPs are the channel through which earth science information originates.

This paper examines the capabilities that various Earth Science Information Partners (ESIPs) may wish to offer in such a federation. We present technologies available to provide these capabilities, with emphasis on components of ECS which may be reused by Earth Science Information Partners.

Keywords: Federation, ESIP, data providers, extended data providers, Earth Science Information Partner, ECS, reuse, evolution, information discovery, high volume production, interoperability

This page intentionally left blank.

Contents

4. Legacy Provider Capabilities

5. Summary

6. References

Figure

Table

1. Introduction

1.1 Purpose

In July 1995, the Board on Sustainable Development Committee on Global Change Research conducted a review of NASA's Mission to Planet Earth program. At the functional nucleus of MPTE is the Earth Observing System Data and Information System (EOSDIS). EOSDIS is a geographically distributed information system which will ingest, process, store, and distribute the data acquired by EOS in-orbit payloads and U.S. observatories. While the EOSDIS review findings were in general favorable, the Committee recommended that NASA

> *reconfigure EOSDIS to transfer responsibility for product generation, publication, and user services to a competitively selected federation of partners in government, academia, and the private sector.*

The committee suggested that such a reconfiguration might accelerate overall system evolvability and help stimulate a wider range of participation from outside the EOS community.

This paper discusses the capabilities that various Earth Science Information Partners (ESIPs) may wish to provide in such a federation and also suggest the means to provide them.

Contributors to this white paper include Lynn Case, Paul Fingerman, Steve Fox, and George Percivall.

1.2 Organization

Section 1 is a brief description of the ECS/Federation evolvability issue brought about by programmatic review of EOSDIS.

The federation, comprised of multiple types of earth science information partners, is further described in Section 2.

Section 3 considers the capabilities new providers may wish to offer and overviews potential technologies available to providers, including those of the EOSDIS Core System (ECS), that offer those capabilities.

In the same vein, Section 4 examines some capabilities which a legacy provider may consider offering.

A summary of our discussion is found in section 5.

1.3 Review and Approval

Being an informal document approved at the Office Manager level, this White Paper does not require formal Government review or approval; however, it is submitted with the intent that review and comments will be forthcoming.

Questions regarding technical information contained within this Paper should be addressed to the following ECS and/or GSFC contacts:

- ECS Contact: Jerry Pisarcik, Sr. Scientist/Engineer, (301) 925-1194, jpisarci@eos.hitc.com
- GSFC Contact: Gail McConaughy, System Architect, (301) 286-7741, gail.mcconaughy@gsfc.nasa.gov

Questions concerning distribution or control of this document should be addressed to:

Data Management Office
The ECS Project Office
Hughes Information Technology Systems
1616 McCormick Drive
Upper Marlboro, MD 20774-5372

2. Earth Science Information Partners (ESIPs)

Earth Science Information Partners (ESIPs) are described by NASA's draft Cooperative Agreement Notices (CAN) as teams or partnerships that will

> *produce and publish/distribute environmental information and provide associated user services in support of Earth System Science. These …ESIPs are expected to demonstrate practical application of new or emerging information technologies and products, including those which are commercially available and to make new or improved information products available to the research community.*

In general, ESIPs provide environmental information and services and have chosen to participate in the federation. They are responsible for the production, archiving, distribution and user services associated with sets of fundamental products. First and foremost ESIPs are information providers. As such, an ESIP's fundamental information or data product could be as simple as a cloud dispersion matrix of the southern Florida peninsula or as complex as a geothermic analysis mapping time-correlated histograms of heat flow across the earth's surface.

Each ESIP is responsible for the local implementation of the federation's system-wide interface layer (SWIL). Using interface protocols and open system standards, the SWIL enables functionality and services throughout the federation. Much like the open system standards adopted by ECS architecture, adherence to federation standards ensures compatibility between federation partners. There are three categories of ESIP, each designated a "type." While there is some overlap, the categorical boundaries are based primarily on provider integration type, responsibility and functionality.

2.1 Type 3 ESIP

Type 3 ESIPs provide data and information products and services to users beyond the Global Change research community. Type 3s are loosely integrated, or autonomous, providers producing and sharing information according to their own schedule, based on their own quality standards, and generally in their own data format and architecture. Loosely integrated providers could range from individual scientists who wish to make their research results widely available for collaborative projects, to stewards of historical data, to providers who are philosophically and architecturally aligned with another federation's vision. Typically, they enter into joint endeavor agreements with institutions like NASA's Mission to Planet Earth (MTPE) in order to extend the benefits of their data products beyond their immediate research community or to enhance the operating capability of EOSDIS. While Type 3 ESIPs may be funded by NASA, other Federal Agencies or by the commercial and non-profit sectors, just as likely, they may be individual researchers operating on separately funded budgets.

2.2 Type 2 ESIP

Type 2 ESIPs are responsible for data and information products and services in support of Earth System Science that are developmental or research-oriented by nature. Emphasizing flexibility and creativity, type 2 ESIPs are key contributors in meeting the federation's advancing research needs. While these ESIPs produce products that may benefit from a less disciplined research environment, their product algorithms require further refinement and maturity; attributes more likely achievable in a loosely integrated "discovery" paradigm. Type 2 ESIPs require mid-to∞ high levels of interoperability and follow the federation data standards closely but not necessarily exactly. Some Type 2 ESIPs might be categorized as 'value-added' information providers. As a value-added provider, the ESIP would provide special products (e.g., tailored subsets, overlays from several sources, special formats) which are of interest specifically to their user community Value-added providers likely have a standard archive and a set of existing interfaces by which users access the data, but would use its association with the federation as a means for its production system.

2.3 Type 1 ESIP

Type 1 ESIPs are responsible for standard data and information products whose production, publishing/distribution, and associated user services require considerable emphasis on reliability and disciplined adherence to schedules. They guarantee data quality, availability, and service levels. Often geographically dispersed, these closely integrated providers support different scientific disciplines, while sharing a common vision of interacting with and providing data to their user community. Environmental information is provided at a pace agreed to with the larger user constituency. Data is archived for long term use and analysis. Type 1 ESIPs support a common data model and architecture enabling their users to readily fuse and analyze distributed, disparate data. They are closely integrated and require the highest level of interoperability. A Type 1 ESIP would provide a full production and distribution capability, including mode management , backup, compression and support of Data Type Services. Today, existing Type 1 ESIP's are Distributed Active Archive Centers (DAAC).

3. New Provider Capabilities

Each potential Earth Science Information Partner brings to the Federation unique products. These products can be standard data processed with special algorithms to yield high resolution and/or better calibrated images. The data may be fused from several instruments yielding information targeting specific study issues such as crop yield prediction or carbon dioxide monitoring. Some of these products already exist and are currently being disseminated to global change research communities; other products are envisioned, logical extensions to current research work.

While a potential ESIP has a unique set of capabilities based on the data or service s/he provides, the user community, and other factors, there are common threads linking each as well.

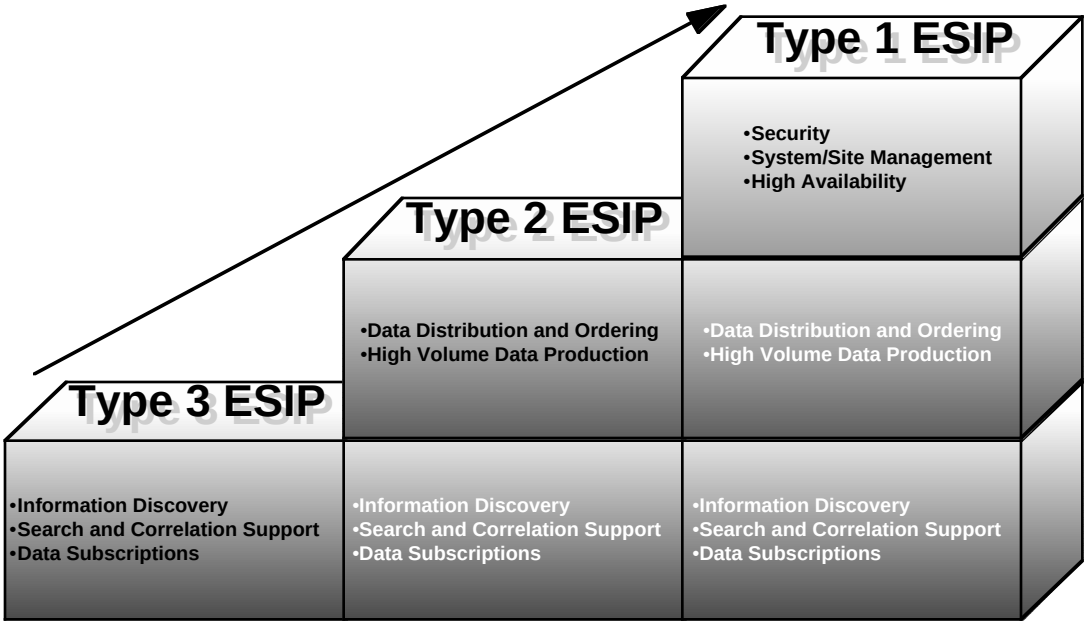

Figure 3-1. ESIP Scenarios. *The provider scenarios build functionality at each step.*

Many providers want to offer their data to a wider research community, but want to do so in a way that minimizes the impact on the resources they offer to user support. Some researchers would like to increase the volume of data provided to users, or the speed at which they can process and disseminate near real time data. Others want to automate the process of distributing data, to free their resources to work on other activities.

The discussions below use these common threads to describe capabilities ESIPs might want to offer. They are "scenarios" illustrating some likely services new ESIPs could offer, as well as candidate implementations an ESIP could choose. Clearly, these scenarios do not exactly fit the particular needs of each ESIP; they are approximations which form the basis for discussion. As Figure 3-1 illustrates, each successive scenario builds functionality and integration, starting with the hypothetical ESIP Type 3 and building to the capabilities of a present-day DAAC.

3.1 Example Capabilities a New ESIP Type 3 Might Provide

The ESIP Type 3 of our scenario is moderately autonomous and wishes to be loosely integrated with the rest of the federation. S/he wishes to make local data holdings available to a wider global change community while maintaining existing ties within his/her specific research community. This provider wishes to offer information that complements EOS data and can be fused with it for interdisciplinary studies. S/he would also like to implement some sort of information management to enable staff at the local site to track data holdings easily. This provider is willing to adopt standards to make information delivery easier, but s/he is not likely to spend very much money integrating new software.

3.1.1 Information Discovery

As with many commercial endeavors, environmental information providers need to reach their user customer market. Traditionally, this market is reached through "advertising" ranging from conference announcements to paid visibility through other services. In today's world, one of the most effective and economical ways to advertise is by means of the Internet and the World Wide Web.

World Wide Web registry services are popular vehicles for alerting the general public that you have something they might need/want. By registering with one or more of the continually increasing number of World Wide Web services (e.g., Yahoo! or Alta Vista), a provider can reach a worldwide audience.

Another approach to advertising data and information services is that invented and used by ECS and CEO. Enabling users access to information relevant to earth science in addition to services like client software, these advertising services specifically and automatically target the earth science research community announcing that new and useful data, tools or services are available from the listing providers. Information contained in the ECS advertising service is entered by the provider or generated at ingest by a data server. In addition to supplying information about a service to other ECS components, the advertising service allows direct querying of information by users through interactive browsing and searching.

3.1.2 Search and Correlation Support

If the provider envisions users correlating EOS data with the information they provide, s/he may wish to enable them to search on the same criteria that they can on EOSDIS core systems. At a minimum, this would enable researchers to select the subset of the provider's information

(e.g., geographic location and time) that is relevant to their study topic, then retrieve the information to their local sites for manual correlation with other data sets. Or the provider may choose to support automated correlation by integrating the data with the ECS distributed information management tools.

Using the Committee on Earth Observing Systems Guidelines for international interoperability, a provider could offer basic search criteria, including:

- Sensor/Platform
- Time/Date
- Location (position)
- Geographical Zone
- Data Quality
- Orbit Characteristics
- Coincident Surface Data
- Instrument Attributes
- Geophysical Attributes.

One readily available data model which supports these and many other search criteria is the EOSDIS data model. The EOSDIS data model is based on analysis of data products for instruments spanning the range of data types, services, and other characteristics of data to be managed by the system. While the data model addresses many issues other than search and correlation (such as optimized storage and browse issues), for purposes of this discussion, we focus on the inventory metadata facilitating such searches and the data dictionary service.

Software supporting access to this data model includes the Science Data Server which implements complex searches and joins of metadata and the Data Management Subsystem described in further detail in the section on Legacy Data Providers. New providers may also use this software to provide the following main functions:

- Provide end-users with a consolidated logical view of a distributed set of data repositories
- Correlate information offered by disparate information providers
- Allow end-users to obtain descriptions for the data offered by these repositories. This also includes descriptions of attributes about the data and the valid values for those attributes.

In addition, the ECS Data Management Subsystem includes a Data Dictionary Service which can be used to translate among differing fields and terminologies. The Data Dictionary stores and provides access to descriptions of data products, their attributes, and the valid values of those attributes. Users query the Data Dictionary to get these descriptions to enhance their knowledge of the system and what it provides. Terms can have different meanings based on the context they are used in. For example, the “Sea Surface Temperature” as a geophysical parameter can have a different meaning (units, location, etc.) dependent on which data product the parameter is used.

The Data Dictionary provides these details to the client. The client is responsible for differentiating the different meanings to the user and making it clear which definition is being used and when.

3.2 Example Capabilities a New ESIP Type 2 Might Provide

The ESIP Type 2 of our scenario has been processing and providing information to colleagues for years as part of an informal information exchange network. This provider applies unique algorithms to EOS-type data to provide higher resolution or better calibrated results that s/he and other researchers use for tracking trends not detectable on standard data products. Lately, this researcher has found that demand for the special data has increased, both in terms of the number of users and the extent of global coverage. S/he would like to scale up the capacity of the data processing center to meet this and future demand for the data, with minimal increase in labor. Our scenario for a new ESIP type 2 provider assumes that s/he will be providing potentially large volumes of data to a global change research community. In addition to using the information discovery and browsing functions of our hypothetical ESIP type 3, this provider will be interested in serving large volumes of scientific data quickly and efficiently to his/her users.

3.2.1 Data Ordering and Distribution

One of the ways a high volume data producer could streamline data delivery is through a suite of data ordering and distribution tools. For instance, after browsing through the provider's inventory to locate appropriate data sets, a researcher may wish to order those data for delivery to his/her local site for closer examination. To facilitate this, a provider could implement a data ordering function which is part of the user interface and analysis tools s/he provides the user community. A data ordering tool which works equally well with all of the tools available to users would be best, as it presents a consistent view to the users and minimizes custom development. Within ECS, this function has been achieved with the Product Request Tool, a component of the ECS Client System Workbench which uses data references obtained with the Earth Science Search Tool, custom search tools, or other references.

User needs, the volume of data to be delivered, network availability, and many other factors influence the means by which data is distributed. Ideally, the data distribution system is flexible enough to accommodate these many factors. Small data volumes can be delivered electronically; usage patterns may determine whether data products are sent directly to the user community or left on a server to be copied at the scientists' convenience. For larger data sets, the provider may wish to deliver data via physical media, including via 4 mm or 8mm tape, 6250 bpi 9-track tape, FAX 3480/3490 or CD-ROM.

Within ECS, Data Distribution (DDIST) orchestrates the delivery of data to its end destination, from enabling user electronic copies to supporting physical media distribution on many possible media types. DDIST also supports operator management of distribution by allowing operators to view, cancel, suspend/resume, and change the priorities of requests.

Through the use of standard data formats, a provider could offer a high level of quality and consistency in their data product(s), as well as maximize portability among researchers. By using the NCSA-developed Hierarchical Data Format (HDF) with EOS conventions, they could also

take advantage of existing, readily available tools to read, write and browse such formatted data. Besides ensuring portability among global change researchers, these tools take advantage of HDF-EOS' self-describing format to exchange highly informative data among sites, and they provide quality control functions.

3.2.2 High Volume Data Production

In order to support searching with additional services such as automated data acquisition and ordering, the Type 2 provider may wish to offer an automated and flexible method for data production capable of ingesting and archiving data on a scheduled or ad hoc basis. Helpful capabilities would include: automated metadata extraction or generation, the ability to schedule whatever processing and loading services are required to archive the data and make it available with its associated metadata, and performance of management functions capable of handling deviations from operational plans. A provider could consider reusing existing technology such as the ECS Data Ingest and Planning and Data Processing (PDPS) subsystems to achieve these capabilities.

The Ingest subsystem deals with the initial reception of all data received at an EOSDIS facility and triggers subsequent archiving and processing of the data. Given the variety of possible data formats and structures, each external interface, and each ad-hoc ingest task may have unique aspects. Therefore, the ingest subsystem is organized into a collection of software components (e.g., ingest management software, translation tools, media handling software) from which those required in a specific situation can be readily configured. The resultant configuration is called an ingest client. Ingest clients operate on a continuous basis serving routine external interfaces; or they may exist only for the duration of a specific ad-hoc ingest task.

The PDPS provides the functions needed to plan routine data processing, schedule on-demand processing, and dispatch and manage processing requests. The subsystem provides access to the data production schedules at each site, and provides management functions for handling deviations from the schedule to operations and science users. To support the complexity of the demands on the processing environment, it was decided to separate long term planning as a custom software implementation from short term planning and scheduling as provided by Autosys, commercial off the shelf software. Production planners will be able to study the impact of changes in resource allocation or processing plans (e.g., the introduction of substantial reprocessing) on overall schedules. At the same time, operations personnel can make decisions regarding the daily workload in a timely fashion, interacting with a highly responsive database, and using a robust commercial product specifically designed for that purpose.

3.2.3 Data Subscriptions

The use of subscription services in which a user would subscribe to receiving data, papers, algorithms/models, and information on new services relevant to their research interest is a potentially powerful tool. Subscriptions are a convenient means for end-users, who have interest in specific scientific disciplines and their related products, to automatically obtain information of special interest as these products become available. Some products may be known future products that are scheduled to be processed or other currently unknown products that may

become available in the future. ECS has implemented a Data Subscription Service which uses a basic event-action interface with which users indicate a desired action when an event of interest occurs. For example, when the event “new generation of CERE03 products available” occurs, the system responds with the predefined automated action “e-mail John Doe that it is available for ordering.”

In a system such as ECS, a subscription to products is based upon a users access privileges and account status. The user submits a subscription to the Subscription Service, describing what data products they would like to see and any supporting documentation that is desired. Actions that may be taken upon receipt of a document that fulfills the request, such as notification of the user, are also specified in the subscription.

3.3 Example Capabilities a New ESIP Type 1 Might Provide

Today, an ESIP Type 1 is a Distributed Active Archive Center (DAAC). In addition to the capabilities the Types 2 and 3 providers offer, a DAAC is responsible for long term storage and distribution of EOS data. As the provider fully responsible for the original data, the DAAC must protect this data from damage, intentional or accidental, through increased security. The DAACs also offer service guarantees to researchers using their computers, implying a high level of system management to ensure those resources are available.

3.3.1 Security and System / Site Management

With the widening scope brought by the federation approach, it is likely that some data sets advertised in the community will have restrictions related to copyright, security, or commercial confidentiality. Requirements for security will vary with the information being provided, its uses, and its user community. Foreknowledge about which security restrictions apply to them allows users to predetermine committing resources in accessing that data. Vulnerabilities to threats are identified and quantified, along with the damage that would result if the threat were successful. A baseline is established to measure the cost-effectiveness of proposed countermeasures. While this is difficult and not cost effective to accomplish on its own, a provider considering offering access restrictions to a given data set because the results need to be validated before a wider research community release would allow the Type 1 ESIP to handle its security.

Within ECS, the security function is performed by the Communication and System Management subsystem. This subsystem supplies system-wide security management and authentication services. Using off-the-shelf technology augmented with some custom software, it provides a communications environment which allows software objects to communicate with each other reliably, synchronously (as well as asynchronously), via interfaces that make the location of a software object and the specifics of the communications mechanisms transparent to the application. In addition, by the Communication and System Management subsystem provides the infrastructural services for the distributed object environment. Based on the Distributed Computing Environment (DCE) from the Open Software Foundation, these services include those needed to develop distributed applications, distributed file services, directory and naming services, security services, and time services.

Security management services are provided enterprise-wide through the Management Services Subsystem (MSS - See System Availability below.), a component of by the Communication and System Management subsystem. The Management Services Subsystem applications provide services such as event fault, performance, security, and accountability management. It manages both ECS' network host and application resources and provides administrative support to the defensive infrastructure employed (e.g., DCE, application-layer gateways, network firewalls) by ECS to defend against a variety of logical threats. It allocates these management services to both the system-wide and local levels. Thus, with few exceptions, management service are fully decentralized, with no single point of failure that would preclude authorized user access, nor permit unauthorized user access.

3.3.2 Guaranteed System Availability

The ECS Management Subsystem provides enterprise management (network and system management) for all ECS resources: commercial hardware (including computers, peripherals, and network routing devices), commercial software, and custom applications. With few exceptions, the management services will be fully decentralized, such that no single point of failure exists which would preclude the system from continuing to operate; or causing system operations and management to come to a halt. ECS selected HP OpenView as the centerpiece of its system management solution, and is augmenting it with other commercially available "agents", as well as custom developed software (e.g., the applications interfaces mentioned above). Type 1 ESIP providers required to meet certain criteria for guaranteed system availability, could easily implement the ECS system availability guidelines. The ECS architecture provides a design for extensible object addition. Enterprise management relies on the collection of information about managed resources, and the ability to send notifications to those resources. For interaction with network devices, computing platforms, and some commercial off the shelf software, MSS relies on software "agents."

While the reliance on COTS software is mandated, a large portion of the ECS applications software is custom developed. Some of this software - the science software - is externally supplied. For these components, the Management Services Subsystem provides a set of interfaces through which these components can provide information to system managers (e.g., about events such as the receipt of a user request or the detection of a software failure). These interfaces also allow applications to accept manual commands from M&O consoles (e.g., an instruction to shut down a particular component). Applications which do not interact with subsystem directly will be monitored by software which acts as their "proxies."

This page intentionally left blank.

4. Legacy Provider Capabilities

As noted above, an existing information provider may choose to join a federation of providers to increase access to his or her data by others as well as gain access to additional data. An information provider in this instance may have many of the capabilities previously discussed and therefore may be more focused on issues like interoperability with the members of the federation. For legacy type providers, the architectural infrastructure is likely already in place. Concentration on services offered would be the crux of the legacy provider's domain. Examples of the capabilities a legacy provider might offer include information discovery, User Administration/Barter/Currency Scheme and Interoperability with Other Providers, and are delineated below.

4.1 Information Discovery

The data and services a provider offers needs to be advertised to users. This capability is described in Section 3.1.1a and is appropriate for a legacy ESIP as well. The provider is able to advertise only those services it wishes to support in the federation; there is no minimum set of services which a site must provide.

4.2 User Administration/Barter/Currency Scheme

The federation cannot assume that there is a single policy regarding user accounting, billing, and authentication. However, the network is likely to include providers which do perform user accounting functions. At one end of the spectrum, each provider could implement his/ her own system for charging user fees, including designing information requests and securing user information from interception. At the other end of the spectrum, a user might interact with a single distributed account management system, maintaining a debit / credit account that is valid with all information providers and ordering data and services through a standard interface.

An ESIP may wish to offer system administration functionality to correlate user's accounts, access rights, etc. across multiple networks. This may include implementation of a bartering or currency exchange mechanism that functions across systems. The federation must offer an API to its administration and accounting systems such that providers can insert their own administrative software if they so desire.

Within ECS, user accounts are the basis upon which services are provided to end-users for subscription and order requests for distribution of products. Information required by subsystem processes is provided through the Accountability Management Service (AMS). Accounts are necessary to provide for billing, distribution of products that require cost reimbursement, accounts receivable and accounts payable processing, and tracking.

4.3 Interoperability with Other Providers

An ESIP may wish to supply tools for interoperating with another provider. A provider's data products, data production algorithms, analysis tools, and computing resources can all be made accessible to the federation by what is referred to in ECS as system interoperability. The ECS architecture offers two types of interoperability that may be considered for reuse among external data providers: publishing of available services and one-stop shopping.

The first type is a loosely coupled interoperability through the publishing of available services. In this type of interoperability the data provider supplies to ECS some information to access services located at his/her site. In most cases, however, the interaction between user and data provider does not require the ECS infrastructure after initial service installation. The ECS Advertising Service provides the data providers with the ability to publish their services.

The three Advertising Service options can be defined as:

1. *Internet Service*: Specifying a particular MIME type, this option allows users searching advertisements in the Advertising Service the ability to connect directly to services through the Web browser currently in use. Installation of a MIME type to the ECS client desktop creates an object on the Desktop which spawns the Web browser opening the appropriate URL and connecting to the service automatically.

2. *Data Provider Supplied Client*: This option allows the data provider to supply a client program that will work with a signature type service. This may be a service resident at the data provider's site or it may be a service resident within ECS. For a service resident at the data provider's site, the ECS infrastructure may not be involved with the client∞ server connection. For ECS services, data providers supply clients, that use the ECS infrastructure, a specialized interface for the ECS service.

3. *ECS Client with Local Data Access Service (LDAS) at Provider Site*: Collectively called the Local Data Access Service (LDAS), the LDAS is a subset of the ECS Data Server, Ingest, and Communications services. It provides local management, search, and access to metadata. In this option, the data provider obtains and configures an LDAS, then advertises the services and exports the schema to the Data Dictionary service. The advertisement which will be a signature type similar to those advertised by ECS' Science Data Servers specifies the ECS client software. Users installing this option get a desktop object which points the client directly to the data provider site.

The second type of interoperability requires a closely coupled interaction between the user, the ECS infrastructure, and the data provider supplied services. In this type of interoperability, the ECS client can issue one search that accesses both ECS sites and the sites of external data providers. The data provider must supply to the Advertising Service both a description of the services provided and a description of the conceptual schema of its holdings to the Data Dictionary Service. This type of interoperability is known as "one-stop shopping."

We define One-Stop shopping as the use of the ECS client when searching and accessing data regardless of the architecture, query language, or database management system at individual sites. ECS provides this capability through its layered architecture which hides the underlying

languages and protocols from the user and client applications in use by the customer. While the Advertising Service provides the capability for data providers to supply services to users, there is a user burden involved in using multiple interfaces to access the data. The one-stop shopping options allow the data provider to appear as if it was one of the ECS DAACs. The user needs just one client to access both ECS DAACs and other data providers.

Table 4-1, ECS Reuse Matrix, identifies the ECS Release B subsystems of interest to each ESIP type.

Table 4-1 ECS Reuse Matrix

Subsystem	Component	ESIP 3	ESIP 2	ESIP 1	Legacy Provider
Client	Desktop		❍	❍	
	Workbench		❍	❍	
Interoperability	Advertiser	●	●	●	
Data Management	Local Information Mgr.	❍	❍	●	
	Distributed Information Mgr.	❍	❍	●	
	Data Dictionary	●	●	●	
Data Server	Science Data Server		❍	❍	
	Document Data Server				
	Storage Mgmt.				
	Data Distribution Service	❍	●		
Ingest	Ingest Services		●	●	
Planning	Production Planning		●	●	
Data Processing	Processing		❍	●	
	Toolkit		❍	●	
Infrastructure	Communication		❍	●	
	Subscription Server	●	●		
System Mgmt.	Management s/w		❍	●	
also:					
Data standards	EOSDIS Data Model	●	●		
	HDF-EOS interchange		●	●	

Key: ● = specifically noted ❍ = partial mention

This page intentionally left blank.

5. Summary

The capabilities a provider offers may be affected by any or all or the following:

1. ESIP Type
2. Level of Interoperability
3. New or Legacy provider

Properly designed and implemented, the federation will provide easy access for all users and providers.

By leveraging the capabilities provided by ECS, a variety of choices exist for providers wishing to make data, algorithms, and analysis tools available to the Earth science community.

This page intentionally left blank.

6. References

1.	Issues to Consider in a Federated Environment	ECS	170-WP-009-001
2.	ECS Application Programming Interface (API) Interface Definition Document (IDD)	ECS	819-RD-001-001
3.	SDPS Database Design and Database Schema Specifications	ECS	311-CD-002-005
4.	CSMS Database Design and Database Schema Specifications	ECS	311-CD-003-005
5.	HDF-EOS Primer for Version-1 EOSDIS	ECS	175-WP-001-001
6.	Thoughts on HDF-EOS Metadata	ECS	170-WP-002-001
7.	Catalogue Interoperability Protocol (CIP) - User Requirements Document (URD)	CEOS	CEOS/WG-D/PTT/CIP-URD
8.	Release-B SDPS/CSMS Design Specification Overview	ECS	305-CD-020-002
9.	A Web Based Science Data Server	ESDIS	http://spsosun.gsfc.nasa.gov//DacinaBox.html
10.	Data Product Analysis: Early Results	ECS	MR9402V1
11.	The Harvest Information Discovery and Access System	U Colo	http://harvest.cs.colorado.edu/

This page intentionally left blank.